

Empirical Analysis on Comparing the Performance of Alpha Miner Algorithm in SQL Query Language and NoSQL Column-Oriented Databases Using Apache Phoenix

Kunal Gupta Astha Sachdev Ashish Sureka
 IIIT Delhi (IIITD)
 New Delhi, India
 ABB Corporate Research
 Bangalore, India

Abstract

Process-Aware Information Systems (PAIS) is an IT system that support business processes and generate large amounts of event logs from the execution of business processes. An event log is represented as a tuple of CaseID, Timestamp, Activity and Actor. Process Mining is a new and emerging field that aims at analyzing the event logs to discover, enhance and improve business processes and check conformance between run time and design time business processes. The large volume of event logs generated are stored in the databases. Relational databases perform well for a certain class of applications. However, there are a certain class of applications for which relational databases are not able to scale. To handle such class of applications, NoSQL database systems emerged. Discovering a process model (workflow model) from event logs is one of the most challenging and important Process Mining task. The α -miner algorithm is one of the first and most widely used Process Discovery technique. Our objective is to investigate which of the databases (Relational or NoSQL) performs better for a Process Discovery application under Process Mining. We implement the α -miner algorithm on relational (row-oriented) and NoSQL (column-oriented) databases in database query languages so that our algorithm is tightly coupled to the database. We present a performance benchmarking and comparison of the α -miner algorithm on row-oriented database and NoSQL column-oriented database so that we can compare which database can efficiently store massive event logs and analyze it in seconds to discover a process model.

Keywords: Apache Hadoop, Apache HBase, Apache Phoenix, Column-Oriented Database, Hadoop Distributed File System (HDFS), MySQL, Process Mining, Row-Oriented Database.

1 Research Motivation and Aim

A PAIS is an IT system that manages and supports business processes. A PAIS generates data from the execution of business processes. The data generated by a PAIS like Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) [20] is in the form of event logs (represented as a tuple $\langle \text{CaseID}, \text{Timestamp}, \text{Activity}, \text{Actor} \rangle$). In an event log, a particular CaseID, that is a process instance, has a set of activities associated with it, ordered by timestamp. Process Mining is new and emerging field which consist of analyzing event logs generated from the execution of business process. The insights obtained from event logs helps the organizations to improve their business processes. There are three major techniques within Process Mining *viz.* Process Discovery, Process Conformance and Process Enhancement [26]. The classification of Process Mining techniques is based on whether there is a priori model and how the a priori model is used, if present. In this paper we focus on Process Discovery aspect of Process Mining. In Process Discovery, there is no a priori model. Process Discovery aims to construct a process model, which is a computationally intensive task, from the the information present in event logs. One of the most fundamental algorithm under Process Discovery is the α -miner algorithm [24] which is used to generate process model from event logs.

Before the year 2000, majority of the organizations used traditional Relational Database Management System (RDBMS) to store the data. Most of the traditional relational databases focus on Online Transaction Processing (OLTP) applications [18] but are not able to perform certain Online Analytical Processing (OLAP) applications efficiently. Row-oriented databases are not able to perform certain analytical functions (like Dense_Rank, Sum, Count, Rank, Top, First, Last and Average) efficiently but work fine when we need to retrieve the entire row or to insert a new record. Recent years have seen the introduction of a number of NoSQL column-oriented database systems [23]. These database systems have been shown to perform more than an order of magnitude better than the traditional relational database systems on analytical workloads [6]. NoSQL column-oriented databases are well suited for analytical queries but result in poor performance for insertion of individual records or retrieving all the fields of a row. Another problem with traditional relational databases is impedance matching [9]. When representation of data in memory and that in database is different, then it is known as impedance matching. This is because in-memory data structures use lists, dictionaries, nested lists while traditional databases store data only in the form of tables and rows. Thus, we need to translate data objects present in the memory to tables and rows and vice-versa. Performing the translation is complex and costly. NoSQL databases on the other hand are schema-less. Records can be inserted at run time without defining any rigid schema. Hence, NoSQL databases do not face the problem of impedance matching.

There are certain class of applications for which row-oriented databases are not able to scale like real time messaging System of Facebook. To handle such class of applications, NoSQL database systems were introduced. Process Discovery is a very important application of Process Mining. Our aim is to examine an approach to implement a Process Discovery α -miner algo-

rithm on a row-oriented database and a NoSQL column-oriented database and to benchmark the performance of the algorithm on both the row-oriented and column-oriented databases. A lot of research has been done in implementing data mining algorithms in database query languages. Previous work suggests that tight coupling of the data mining algorithms to the database systems improves the performance of the algorithms significantly [13]. We aim to implement α -miner algorithm in Structured Query Language (SQL) so that our Process Discovery application is tightly coupled to the database.

Combination of Hadoop¹ component and NoSQL column-oriented databases allow accessing large data efficiently and storing data easily as compared to single machine databases [25]. There are various NoSQL column-oriented databases [23] but for our current work, we will focus on Apache HBase² (NoSQL column-oriented database) and MySQL³ (row-oriented database). To perform analytical functions, NoSQL column-oriented databases either use MapReduce programming model or use their own simple query language that just supports create, read, update and delete (CRUD). They do not support SQL interface. We integrate Apache Phoenix⁴ (SQL layer over HBase) into HBase to support SQL interface in it. It converts SQL queries to HBase scans rather than MapReduce jobs. It executes converted scans in parallel over the regions in a regionserver and targets low latency query over HBase tables as compared to MapReduce framework and client API's.

Main research aim presented in this paper is-

1. To investigate an approach to implement α -miner algorithm in SQL. The underlying row-oriented database for implementation is MySQL using InnoDB⁵ engine.
2. To investigate an approach to implement α -miner algorithm on column-oriented database HBase using Phoenix and HDFS.
3. To conduct a series of experiment on publicly available real world dataset, to compare the performance of α -miner algorithm on both the databases. The experiment considers multiple aspects such as α -miner stepwise execution, bulk loading across various datasets, write intensive time, read intensive time, disk space of tables, disk space of tables using compression technique, α -miner stepwise execution using compression technique, real time batch wise insertion and real time single record insertion.

2 Related Work

In this Section, we review closely related work to the study presented in this paper and list the novel contributions of our work in context to existing work. We divide related work into

¹<http://hadoop.apache.org>

²www.hbase.apache.org

³<http://www.mysql.com/>

⁴<http://www.phoenix.apache.org>

⁵<http://dev.mysql.com/doc/refman/5.5/en/innodb-storage-engine.html>

following three lines of research:

2.1 Implementation of Mining Algorithms in Row-Oriented Databases

Ordonez et al. investigate an approach to efficiently implement the EM algorithm in SQL[4]. They perform clustering of large datasets. They effectively handle high dimensional data, a high number of clusters and more importantly, a very large number of data records. Xuequn Shang et al. present an efficient implementation of frequent pattern mining in relational databases [22]. They propose a concept called Projection Pattern Discovery (Propad). Propad fundamentally differs from Apriori like candidate set generation-and-test approach. This approach successively projects the transaction table into frequent itemsets to avoid making multiple passes over the large original transaction table and generating a huge set of candidates.

2.2 Implementation of Mining Algorithms in Graph Databases, Parallelization and Utility Based Approach

Joishi et al. implement Similar-Task algorithm on relational and NoSQL (graph oriented) databases using only query language constructs [12]. They conduct empirical analysis on a large real world data set to compare the performance of row-oriented database and NoSQL graph-oriented database [12]. Kundra et al. investigate the application of parallelization on Alpha Miner algorithm [14]. They use Graphics Processor Unit (GPU) to run computationally intensive parts of Alpha Miner algorithm in parallel [14]. Anand et al. propose a Utility-Based Fuzzy Miner UBFM algorithm to efficiently mine a process model driven by a utility threshold [1].

2.3 Implementation of Mining Algorithms in Column-Oriented Databases

Mehta et al. conducted a study of data mining algorithms on column-oriented database systems [19]. They study the architecture of open source column-oriented databases and implemented tree based classification algorithm on various column-oriented databases like MonetDB and Infobright. Suresh L et al. presented an implementation of k-means clustering algorithm on column-oriented databases [16]. They introduce an algorithm known as Novel Seeding Algorithm to implement k-means in column-oriented databases. This algorithm identifies the median gaps in the data in each of the columns and using these gaps to identify other clusters by using the difference in the median gaps.

2.4 Performance Comparison of Mining Algorithms in Row-Oriented and Column-Oriented Databases

Hasso conducted common database approach for OLTP and OLAP using an in-memory column database [18]. He presented a comparison of OLAP and OLTP considering row-oriented

database and column-oriented database. Rana et al. implement Apriori algorithm on MonetDB and Oracle database and compare their performance in terms of execution time [7].

2.5 Novel Contribution

In context of existing work, this study presented here makes the following novel contributions. The work presented in this paper is extension of the work presented in [21][10]. The study presented in this paper has several more results which are not present in [21][10] due to limited space in the conference paper.

1. While there has been work done on implementing data mining algorithms on row-oriented databases, we are the first to implement Process Mining α -miner algorithm on MySQL using InnoDB storage engine.
2. While there has been work done on implementing data mining algorithms on column-oriented databases, we are the first to implement Process Mining α -miner algorithm in HBase using Phoenix and HDFS.
3. We present a performance benchmarking and comparison of α -miner algorithm on both MySQL and HBase. We consider multiple aspects such as α -miner stepwise execution, bulk loading across various datasets, write intensive time, read intensive time, disk space of tables, disk space of tables using compression technique, α -miner stepwise execution using compression technique, real time batch wise insertion and real time single record insertion.

3 α -Miner Algorithm

The α -miner algorithm is an algorithm used in discovering Process Mining [24]. It was first put forward by van der Aalst, Weijter and Maruster. Input for the α -miner algorithm is an event log L and output is a process model. The α -miner algorithm consists of scanning the event logs for discovering causality between the activities present in the event log. The basic ordering relations determined by α -miner algorithm are the following:

1. $a \succ_L b$ iff a directly precedes b in some trace. Where a and b can be set of activities or an activity and this relation represents causality.
2. $a \rightarrow_L b$ iff $a \succ_L b \wedge b \succ_L a$.
3. $a \parallel b$ iff $a \succ_L b$ and $b \succ_L a$ in some trace.
4. $a \# b$ iff $a \succ_L b \wedge b \succ_L a$.

Let L be an event log over T , where T is the set of distinct activities present in the event log and σ is a trace in the event log. $\alpha(L)$ is defined as follows.

The stepwise description of the α -miner algorithm can be given as:

1. Step 1 computes T_L (Total Events) which represents the set of distinct activities present in the event log L .
2. Step 2 computes T_I (Initial Events) which represents the set of all the initial activities of corresponding trace.
3. Step 3 computes T_O (Final Events) which represents the set of distinct activities which appear at the end of some trace in the event log.
4. In order to compute Step 4, we compute the relationships between all the activities in T_L . This computation is presented in the form of a footprint matrix and is called pre-processing in α -miner algorithm. Using the footprint matrix we compute pairs of sets of activities such that all activities in the same set are not connected to each other while every activity in first set has causality relationship to every other activity in the second set.
5. Step 5 keeps only the maximal pairs of sets generated in the fourth step, eliminating the non-maximal ones.
6. Step 6 adds the input place which is the source place and the output place which is the sink place in addition to all the places obtained in the fifth step.
7. Step 7 is the final step of the α -miner algorithm that presents all the places including the input and output places and all the input and output transitions from the places.

4 Implementation of α -Miner Algorithm in SQL on Row-Oriented Database (MySQL)

We present a few segments of our implementation due to limited space in the paper. The entire code and implementation can be downloaded from our website⁶. Before implementing α -miner algorithm, we do pre-processing in JAVA to create two tables *viz.* causality table (consist of two column eventA and eventB) and notconnected table (consist of two column eventA and eventB).

1. We create a table eventlog using create table⁷ keyword consisting of 5 columns (CaseID, Timestamp, Status, Activity and Actor) each of which are varchar datatype except Timestamp which is of timestamp datatype. The primary key is a composite primary key consisting of CaseID, Timestamp and Status.
2. We load the data into table eventlog using LOAD DATA INFILE⁸ command.

⁶<https://dl.dropboxusercontent.com/u/48972351/Programing-Alpha-miner-in-MySQL-and-HBase.zip>

⁷<http://dev.mysql.com/doc/refman/5.1/en/create-table.html>

⁸<http://dev.mysql.com/doc/refman/5.1/en/load-data.html>

3. For Step 1, we create a table `totalEvent` that contains a single column (`event`) which is of `varchar` datatype. To populate the table we select distinct activities from the table `eventlog`.
4. For Step 2, we create a table `initialEvent` that contains a single column (`initial`) which is of `varchar` datatype. To populate the table
 - (a) We first select the minimum value of `Timestamp` from table `eventlog` by grouping `CaseID`.
 - (b) Then we select distinct activities from table `eventlog` for every distinct value of `CaseID` where `Timestamp` is the minimum `Timestamp`.
5. For Step 3, we create a table `finalEvent` that contains a single column (`final`) which is of `varchar` datatype. To populate the table
 - (a) We first select maximum `Timestamp` from a table `eventlog` by grouping `CaseID`.
 - (b) Then we select distinct activities from a table `eventlog` for every distinct value of `CaseID` where `Timestamp` is the maximum `Timestamp`.
6. For Step 4, we create five tables *viz.* `SafeEventA`, `SafeEventB`, `EventA`, `EventB` and `XL`. All the five tables contain two columns (`setA` and `setB`) which are of `varchar` datatype.
 - (a) In table `causality` we use `group_concat`⁹ to combine the values of column `eventB` of corresponding value of a column `eventA` and insert the results in a table `EventA`.
 - (b) In table `causality` we use `group_concat` to combine the values of column `eventA` of corresponding value of a column `eventB` and insert the results in the table `EventB`.
 - (c) To populate tables `SafeEventA` and `SafeEventB`-
 - i. Select `setA` and `setB` from tables `EventA` and `EventB`
 - ii. For every value of `setB` in table `EventA`, if value is present in table `notconnected`, insert the corresponding value of `setA` and `setB` in table `SafeEventA`. Repeat the same step for populating table `SafeEventB`.
 - (d) To populate table `XL`, we insert all the rows from the three tables `SafeEventA`, `SafeEventB` and `causality`.
7. For Step 5, we create three tables *viz.* `eventASafe`, `eventBSafe` and `YL`. All the three tables contain two columns (`setA` and `setB`) which are of `varchar` datatype.
 - (a) We create a stored procedure to split the values of column `setB` of table `SafeEventA` on comma separator. Insert the results in `safeA` table.
 - (b) We create a stored procedure to split the values of column `setA` of table `SafeEventB` on comma separator. Insert the results in `safeB` table.

⁹http://dev.mysql.com/doc/refman/5.0/en/group-by-functions.html#function_group_concat

- (c) To populate table eventASafe, insert all the rows from table safeA.
 - (d) To populate table eventBSafe, insert all the rows from table safeB.
 - (e) To populate table YL, insert all the rows from tables SafeEventA, SafeEventB, eventASafe, eventBSafe and causality.
8. For Step 6, we create two tables *viz.* terminalPlace that contains a single column (event) which is of varchar datatype and PL which also contains a single column (Place) which is of varchar datatype.
- (a) To populate table terminalPlace, insert 'i' and 'o' in the table.
 - (b) To populate table PL, we use concat_ws¹⁰ to combine the values of column setA and column setB of a table YL using & separator and insert the results in table PL. Furthermore, we insert all the rows of table terminalPlace into table PL.
9. For Step 7, we create 3 tables *viz.* Place1 and Place2 which consist of two columns (id and value) which are of varchar datatype and FL which consists of two columns (firstplace and secondplace) which are of varchar datatype.
- (a) To populate table Place1, we use concat_ws to combine the values of column setA and column setB of table YL using & separator. Insert the results in column setB of table Place1. Insert all the values of column setA of table YL into column setA of table Place1.
 - (b) To populate table Place2, we use concat_ws to combine the values of column setA and column setB of table YL using & separator. Insert the results in column setA of table Place2. Insert all the values of column setB of table YL in column setB of table Place2.
 - (c) We create a stored procedure to split column setB of table Place1 on comma separator. In stored procedure we create table temp_place2 to insert the results.
 - (d) We create a stored procedure to split column setA of a table Place2 on comma separator. In stored procedure we create table temp_place2 to insert the results.
 - (e) To populate a table FL, insert all the rows from tables temp_place1 and temp_place2. Insert the results of cross join of two tables *viz.* terminalPlace and intialEvent and of table finalEvent and table terminalPlace.

5 Implementation of α -Miner Algorithm on NoSQL Column-Oriented Database (HBase) Using Apache Phoenix

Before implementing α -miner algorithm, we do pre-processing in JAVA to create two tables *viz.* causality table (consist of two column eventA and eventB) and notconnected table (consist of

¹⁰http://dev.mysql.com/doc/refman/5.0/en/string-functions.html#function_concat_ws

two column eventA and eventB).

1. To create table eventlog (Refer section 4 point 1). To load the data in table eventlog, we use MapReduce framework¹¹.
2. For Step 1, 2 and 3 (Refer section 4 point 3, point 4 and point 5)
3. For Step 4, we create three tables *viz.* SafeEventA, SafeEventB and XL. All the three tables consist of two columns (setA and setB) which are of varchar datatype.

Algorithm 1: Populating table SafeEventA and SafeEventB

```

1 Select eventA, eventB from causality.
2 Select setA, setB from notconnected.
3 Compare Function that compares whether set of activities is notconnected.
4 foreach eventA in the causality do
5   | Form single group say grp of all activity present in column eventB. Pass grp to Compare
(a) | function. For any such combination returning true, insert eventA in setA and that
    | combination into setB of table SafeEventA.
6 end
7 foreach eventB in the causality do
8   | Form single group say grp of all activity present in column eventA. Pass grp to Compare
    | function. For any such combination returning true, insert that combination into setA and
    | eventB in setB of table SafeEventA.
9 end

```

(b) To populate table XL, we insert all the rows from three tables SafeEventA, SafeEventB and causality.

4. For Step 5, we create three tables *viz.* EventA, EventB and YL. All the three tables consist of two columns (setA and setB) which are of varchar datatype.

Algorithm 2: Populate table EventA and EventB

```

1 Select setA, setB from SafeEventA.
2 Select setA, setB from SafeEventB.
3 foreach setA,setB in the SafeEventA do
(a) 4 | Delimit value of setB. For all such value setBi, insert setA and setBi in table EventA.
    5 end
6 foreach setA,setB in the SafeEventB do
7   | Delimit value of setA. For all such value setAi, insert setAi and setB in table EventB.
8 end

```

(b) To populate table YL, we insert all the rows from three tables EventA, EventB and causality.

5. For Step 6 (Refer section 4 point 8).
6. For Step 7, we create table FL that consists of two columns (Place1 and Place2) which are of varchar datatype.

¹¹http://phoenix.apache.org/bulk_dataload.html

Algorithm 3: Populating Table FL

```
1 Select setA, setB from YL.
2 Select final from FinalEvents.
3 Select initial from InitialEvents.
4 foreach final in the FinalEvents do
5   | Insert final in column Place1 and 'o' in column Place2 of table FL
6 end
7 foreach initial in the InitialEvents do
(a) 8   | Insert 'i' in column Place1 and initial in column Place2 of table FL
9 end
10 foreach setA,setB in the YL do
11   | If value of column setA has set of activities instead of single activity then delimit. Each
    | split value will be stored in column Place1 and combination of setA and setB in
    | column Place2 of table FL
12   | else choose column setB and delimit. Each split value will be stored in column
    | Place2 and combination of setA and setB in column Place1 of table FL
13 end
```

6 Experimental Dataset

We conduct our study on a publicly available large real world dataset downloaded from Business Process Intelligence 2014 (BPI 2014)¹². The dataset is provided by Robobank Information and Communication and Technology (ICT). The data is related to Information Technology Infrastructure Library (ITIL) process implemented in the bank. ITIL is a process which starts when a client reports an issue regarding disruption of ICT service to Service Desk Agent (SDA). SDA records the complete information about the problem in an Interaction record. If the issue does not get resolved on first contact then an Incident record is created for the corresponding Interaction else the issue is closed. If an issue appears frequently then a request for change is initiated. Robobank provides 4 files in CSV format *viz.* Change records, Incident records, Interaction records and Incident activity records. We import Incident activity records CSV file in MySQL and HBase for benchmarking and performance comparison of α -miner algorithm. Incident activity records file contains 4,66,738 number of records and contains the following fields *viz.* Incident ID, DateTimeStamp, IncidentActivity_number, IncidentActivity_Type, Interaction ID, Assignment Group and KM Number. Out of these we use the following fields:

1. Incident ID: The unique ID of a record in the Service Management tool. It is represented as CaseID in our data model.
2. DateTimeStamp: Date and time when a specific activity starts. It is represented as timestamp in our data model.
3. IncidentActivity_Type: Identifies which type of an activity takes place.
4. Assignment Group: The team responsible for an activity.

¹²<http://www.win.tue.nl/bpi/2014/start>

7 Benchmarking and Performance Comparison

Our benchmarking system consists of Intel Core i3 2.20 GHz processor, 4 GB Random Access Memory (RAM), 500 GB Hard Disk Drive (HDD), Operating System (OS) is Linux Ubuntu 14.04 LTS and Cache of 3 MB. The experiments were conducted on MySQL 5.6 (row-oriented database) and HBase 0.96.1 (NoSQL column-oriented database) with HDFS 2.3.0 as the file system below it and a layer of Phoenix 4.2.1 above it. We conduct series of experiments on a single machine.

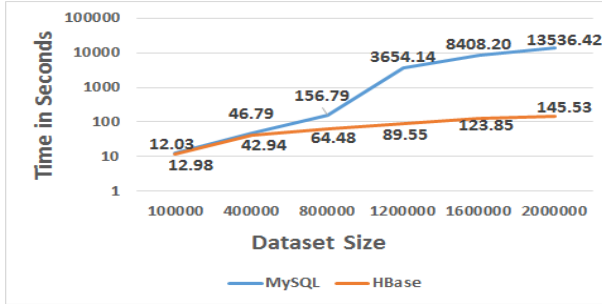
The α -miner algorithm interacts with the database. The underlying data model for implementing α -miner algorithm consists of 5 columns (CaseID, Timestamp, Status, Activity and Actor) each of which are varchar datatype except Timestamp which is of timestamp datatype. The primary key is a composite primary key consisting of CaseID, Timestamp and Status. We use the same data model while performing bulk loading of datasets through the database loader. We take each reading five times for all the experiments and the average of each reading is reported in the paper.

Table 1: Dataset Load Time

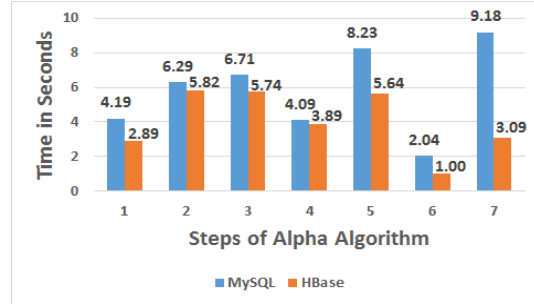
Dataset Size	Load Time in Seconds	
	MySQL	HBase
1,00,000	12.98	12.03
4,00,000	46.79	42.94
8,00,000	156.79	64.48
12,00,000	3654.14	89.55
16,00,000	8408.20	123.85
20,00,000	13536.42	145.53

Table 2: Stepwise Execution Time

Steps	Execution Time in Seconds	
	MySQL	HBase
1	4.19	2.89
2	6.29	5.82
3	6.71	5.74
4	4.09	3.89
5	8.23	5.64
6	2.04	1.00
7	9.18	3.09



(a) Dataset Load Time in Seconds



(b) α -miner Stepwise Execution Time in Seconds

Figure 1: Dataset Load Time and α -miner Stepwise Execution Time

Our first experiment consists of investigating the time taken to perform bulk loading in both the databases across various dataset sizes. Table 1 shows that the average time taken to load data in HBase is 29 times lower as compared to MySQL. Bulk loading in HBase is done using the MapReduce framework. Phoenix has an inbuilt script of MapReduce using which we conduct our experiment. We use two mappers and two reducers for running MapReduce

jobs. The script requires two parameters before running MapReduce *viz.* input file and output file. Input file must be present in HDFS and script creates empty output file in HDFS after executing it. Due to parallelism, all the key-values of the input file are mapped to two mappers and the output of each mapper is passed to two reducers. MapReduce converts all the data of the input file into the format of HFiles (HBase file format) before it handovers to HBase. HFile stores data in key-value pairs and reducers also generate output in key-value pairs. The output of reducers can be stored on multiple HFiles directly without interacting with HBase. At the end all the created HFiles will be handovered to HBase to store on HDFS. Bulk loading in MySQL is done using LOAD DATA INFILE command which is designed for mass loading of records in a single operation as it overcomes the overhead of parsing and flushing batch of inserts stored in a buffer to MySQL server. LOAD DATA INFILE command also creates an index and checks uniqueness while inserting records in a table. Therefore, in case of MySQL, while inserting large datasets, most of the time is spent in checking uniqueness and creating indexes. Fig. 1(a) reveal that when the dataset size increases then the difference between the time taken in loading data in MySQL and HBase also increases. The performance of HBase is better as compared to MySQL because the percentage increase of time in MySQL is 3.5 times more as compared to HBase.

α -miner algorithm is a seven step algorithm (Refer to Section 3.2). Few steps are read intensive (Steps 1,2,3) while few steps are write intensive (Steps 4,5,6,7). We perform an experiment to compute the α -miner algorithm execution time of each step in both MySQL and HBase to examine which database performs better for each step. In MySQL default size of innodb_buffer_pool_size is 8 MB that is used to Cache data and indexes of its tables. The larger we set this value, the lesser is the disk I/O needed to access the data in tables. Table 2 and Fig. 1(b) reveal that the the stepwise time taken in HBase is always lower as compared to MySQL for all the Steps. We conjecture the reason for HBase performing better than MySQL can be the difference in the internal architecture of MYSQL and HBase. For the first three steps, both MySQL and HBase perform full table scans. In case of MySQL, the entire row is first retrieved sequentially and then the specific attributes are retrieved. However, in case of HBase, table is stored on multiple regions and Phoenix performs parallelism on multiple regions of a table leading to better performance of HBase in comparison to MySQL. Furthermore, in HBase, only the specific attributes specified in the query are retrieved. The overhead of retrieving the entire row is not present in HBase. Hence, HBase gives a better performance for the first three steps.

The remaining steps read data from the tables obtained in the first three steps and write it to the tables created during their execution. In MySQL, in order to read the data from a table we need to scan the B-Tree index to find the location of block where data is stored. In case of HBase data is read from the memstore. If values are not in memstore they are read from HDFS. Thus, the read performance of HBase is better as compared to MySQL. Similarly, in MySQL, in order to write data, the entire B-Tree index needs to be scanned to locate the block where

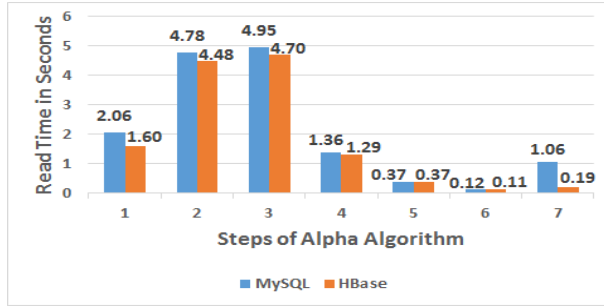
we need to write data. HBase follows log structure merge tree index. In case of HBase, values are written in append only mode. The writes in HBase are sequential because first it writes to WAL (Write Ahead Log) of regionserver and then to memstore of corresponding region. HBase lags in persisting data to disk. Hence, HBase gives better write performance as compared to MySQL. Therefore, the total execution time of α -miner algorithm in HBase is 1.44 times lower than that of MySQL.

Table 3: Read Intensive Time

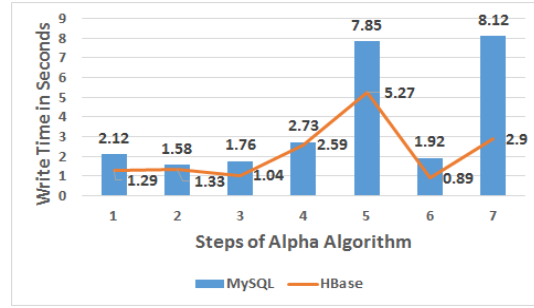
Steps	Read Time in Seconds	
	MySQL	HBase
1	2.06	1.60
2	4.78	4.48
3	4.95	4.70
4	1.36	1.29
5	0.37	0.37
6	0.12	0.11
7	1.06	0.19

Table 4: Write Intensive Time

Steps	Write Time in Seconds	
	MySQL	HBase
1	2.12	1.29
2	1.50	1.33
3	1.76	1.04
4	2.73	2.59
5	7.85	5.27
6	1.92	0.89
7	8.12	2.90



(a) Read Intensive Time in Seconds



(b) Write Intensive Time in Seconds

Figure 2: Read and Write Intensive Time

α -miner algorithm consists of seven steps (Refer to Section 3.2). Few steps are more intensive for read operations while few steps are more intensive for write operations. We conduct an experiment to compare which of the database performs better for read and write operations in both MySQL and HBase. As can be seen from Fig. 2(a) and Table 3, HBase gives better read performance as compared to MySQL for all the Steps. According to us, the reason for HBase giving better read performance can be the difference in the data structure of both the databases. In MySQL, B-Tree index needs to be scanned to find the location of block where the data is stored. In case of HBase data is read as described below-

1. To find the data, HBase client will hit the memstore first.
2. When the memstore fails, HBase client will hit the BlockCache [5].
3. If both the memstore and BlockCache fail, HBase client will locate the target HFiles in HDFS (contains target data) using log structure merge tree and load it into the memory.

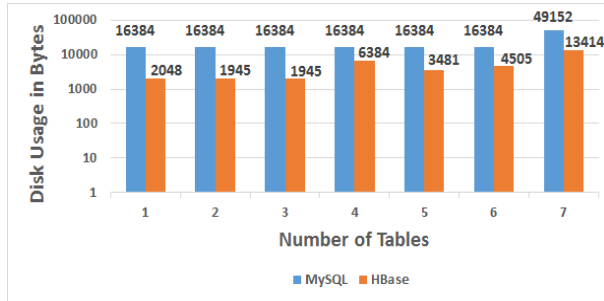
The total time taken to read the data in each of the Step of α -miner algorithm is 1.16 times lower in HBase as compared to MySQL. Fig. 2(b) and Table 4 show that the write performance of HBase is better as compared to MySQL. We believe the reason for HBase giving better write performance can be the difference in the way writes are performed in both the databases. In MySQL, the B-Tree index needs to be scanned to find the location of block where the data needs to be written. Almost all the leaf blocks of B-Tree are stored on the disk. Hence, at least one I/O operation is required to retrieve the target block in memory. Fig. 2(b) illustrates that Step 5 and Step 7 of α -miner algorithm in MySQL are more write intensive than the other steps. We believe the reason can be the generation of maximal sets and places by stored procedures in MySQL. A large number of insert operations are executed in the stored procedure to generate the maximal sets. In HBase we perform the same steps using Java because SQL interface over HBase does not support advanced features of SQL. Writes in HBase are performed by first locating regionserver from zookeeper¹³, then regionserver writes to WAL and finally to memstore of the corresponding region. Phoenix allows to perform parallelism in reading and writing the data on multiple regions of a table stored in HBase regionserver in comparison to sequential reads and writes of MySQL. The total time taken in writing the data in each of the Step of α -miner algorithm is 1.70 times lower in HBase as compared to MySQL. Thus, writes in HBase are more optimized as compared to that in MySQL.

Table 5: Disk Usage of Tables

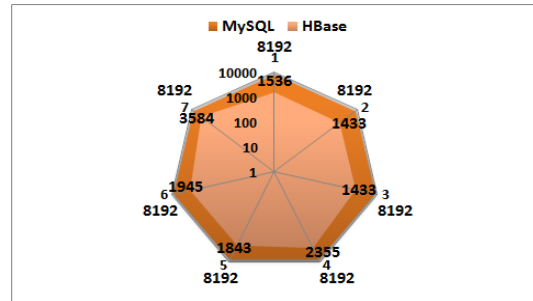
Step wise Tables	Disk Usage in Bytes	
	MySQL	HBase
Step 1	16384	2048
Step 2	16384	1945
Step 3	16384	1945
Step 4	16384	6348
Step 5	16384	3481
Step 6	16384	4505
Step 7	49152	13414

Table 6: Disk Usage of Tables With Compression

Step wise Tables	Disk Usage in Bytes	
	MySQL	HBase
Step 1	8192	1536
Step 2	8192	1433
Step 3	8192	1433
Step 4	8192	2355
Step 5	8192	1843
Step 6	8192	1945
Step 7	8192	3584



(a) Disk Usage of Tables in Bytes



(b) Disk Usage of Tables in Bytes with Compression

Figure 3: Disk Usage of Tables in Bytes with and without Compression

¹³<http://www.zookeeper.apache.org>

We perform an experiment to investigate which database can efficiently store results of each Step of α -miner algorithm in tables with minimum disk space. Table 5 and Fig. 3(a) reveal the disk space occupied by tables created in each step of α -miner algorithm. We include only data length (excluding the size of index tables) in disk space of table because we did not create index for any of the tables. Experimental results show that HBase on an average uses disk space 6 times lower than MySQL for tables created at each step of the algorithm. Hence, cumulative disk space for storing all the tables in MySQL is 147456 bytes while for HBase is 33722 bytes. We believe the underlying reason for MySQL occupying more space is the difference in the way memory is allocated to tables in both the databases. In MySQL, the operating system allocates fixed size blocks of size 16 KB for the data to be stored in a table. Number of blocks assigned to a table is computed by dividing the dataset size by the block size. In MySQL if set of blocks or one block has been allocated for a table then that set of blocks or block can be used only by that table. Either data in a table completely utilizes the space of all blocks or the space of the last block is unutilized. Storing smaller size file (< 16 KB) in 16 KB block leads to under utilization of space and the remaining space cannot be utilized by other files.

HFile is a file format of HBase which is stored over HDFS block (default size is 64 MB). Maximum size of a HFile is 64 KB after which a new HFile needs to be created. HFiles are created when memstore reaches its threshold value (default value is 64 MB) or commit occurs. When memstore reaches its threshold value it flushes 64 MB data of key-value pairs and creates 1024 numbers of HFiles. If commit occurs before it reaches the threshold value then it flushes only that amount of data present in a memstore. HFile size will be equivalent to flushed amount of data from memstore. HDFS allocates blocks for incoming files by dividing the file size with the block size. For example, we have a system with 300 MB HDFS block size. To store a 1100 MB file, HDFS will break that file into three 300 MB blocks and one 200 MB block size and store it on the datanodes. The 200 MB file is not exactly divisible by 300. Therefore, the final block of the file is sized as modulo of the file size by block size, i.e a 200 MB block size. Similarly, the same process is applied to the HFiles of HBase for storing in HDFS. We conclude that the disk space for each table created in each step is more efficiently utilized in HBase as compared to MySQL.

A way to utilize disk space efficiently is by using the well known compression technique. Data compression enables smaller database size, reduced I/O and improved throughput. We conduct an experiment to compute the disk space occupied by tables at each Step of the α -miner algorithm using compression technique. When we compare the disk space occupied by each table without compression and with compression technique we observe that the compression ratio (Actual size of table/Compressed size of table) is better in MySQL as compared to HBase. As can be seen from Table 5 and Table 6, the compression ratio in MySQL for Step 7 is equal to 6 (49152/8192) while the compression ratio in HBase for Step 7 is equal to 3.7. Minimum and maximum compression ratio in HBase is 1.3 and 3.7 respectively while in MySQL is 2 and 6 respectively. We believe the reason for MySQL having a higher compression ratio can

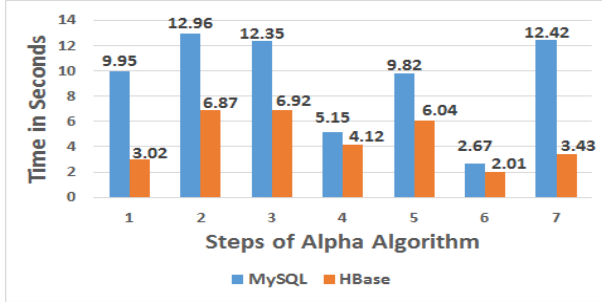
be the difference in the compression techniques used by both the databases. MySQL uses the zlib compression technique which provides a better compaction using only six bytes of header and trailer of compressed block. HBase uses gzip compression technique and gzip wrapper uses a minimum of eighteen bytes of header and trailer for compressed block. The maximum compression ratio provided by MySQL is 2 times more as compared to HBase. In the context of α -miner algorithm, MySQL performs better than HBase in utilizing the disk space when compression technique is applied.

Table 7: Stepwise Execution Time with Compression

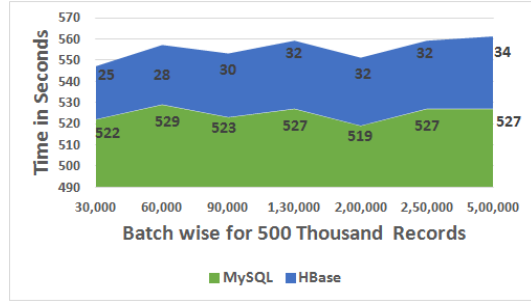
Steps	Execution Time in Seconds	
	MySQL	HBase
1	9.95	3.02
2	12.96	6.87
3	12.35	6.92
4	5.15	4.12
5	9.82	6.04
6	2.62	2.01
7	12.42	3.43

Table 8: Batch wise Insertion Time

Batch Size for 500 Thousand Records	Batch wise Insertion Time in Seconds	
	MySQL	HBase
30,000	522	25
60,000	529	28
90,000	523	30
1,30,000	527	32
2,00,000	519	32
2,50,000	527	32
5,00,000	527	34



(a) α -miner Stepwise Execution Time with Compression



(b) Batch wise Insertion Time in Seconds

Figure 4: α -miner Stepwise Execution Time with Compression and Batch wise Insertion Time in Seconds

We conduct an experiment to examine the time taken by each Step of α -miner algorithm with compression technique. In α -miner algorithm we create tables in each Step with the compression keyword. Table 7 and Fig. 4(a) illustrate that the performance of HBase is better as compared to that of MySQL for each Step of α -miner algorithm. We believe the reason for HBase giving better step wise execution time, with compression enabled can be the difference in the way compression is performed in both the databases. MySQL uses a block size of 1 KB, 2 KB, 4 KB, 8 KB and 16 KB. The default block size after compression in MySQL is 8 KB. Suppose the size of the compressed block is 5 KB. The block will then be uncompressed, split into two blocks and then recompressed into blocks of size 4 KB and 1 KB. All the data in a table is stored in blocks comprising a B-Tree index. The compression

of B-Tree blocks is handled differently because they are frequently updated. It is important to minimize the number of times B-Tree blocks are split, uncompressed and recompressed. MySQL maintains system information in B-Tree block in uncompressed form for certain in-place updates. MySQL avoids unnecessary uncompression and recompression of blocks when they are changed because it causes latency and degrades the performance. HBase does not have fixed block size constraint after compressing the block. We conjecture that another reason for HBase giving a better stepwise execution time, with compression enabled can be the difference in the internal architecture of both the databases that was explained in experiment (Refer to Table 2 and Fig. 1(b)). From Table 2 and Table 7, we infer that the total execution time of α -miner algorithm in MySQL is 2 times more as compared to HBase using compression technique. We compare the total time taken in executing α -miner algorithm without compression and with compression technique in MySQL and HBase. We observe that total time taken in executing α -miner algorithm by HBase without compression technique is 1.33 times lower than HBase with compression technique. Similarly, MySQL without compression technique is 1.60 times lower than MySQL with compression technique.

In all the experiments described above the event logs generated from business processes is stored in a CSV file and then loaded in the database. In the context of Process Mining, PAIS are getting continuously updated with event logs. We setup our experiment to import the event logs directly into the database server from a client application, that is real time data (event logs) loading. The real time loading experiment can be conducted in two ways *viz.* batch insertion and single row insertion. In the batch insertion, the client application inserts 5,00,000 records in different batch sizes. The results of batch insertion are shown in Fig. 4(b) and Table 8. We believe that batch insertion might be faster than single record insertion because when we execute a batch, then multiple records in a batch are inserted in a table in a single round trip.

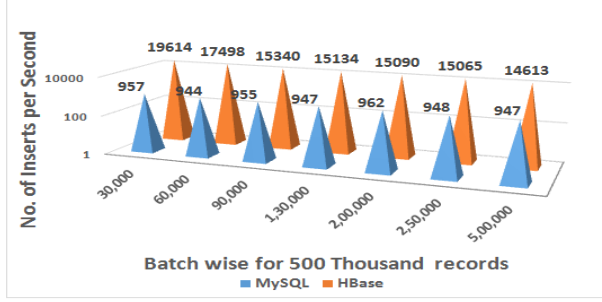
Table 9: Number of Inserts per Second in Batch

Batch wise for 500 Thousand Records	Number of Inserts per Second	
	MySQL	HBase
30,000	957	19614
60,000	944	17498
90,000	955	15340
1,30,000	947	15134
2,00,000	962	15090
2,50,000	948	15065
5,00,000	947	14613

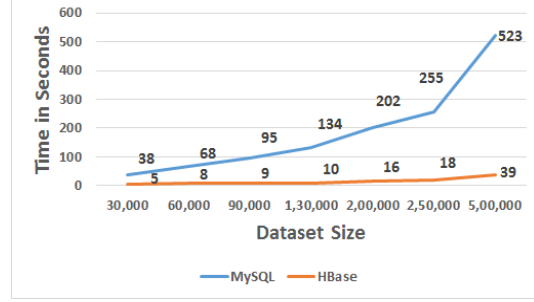
Table 10: Single Row Insertion Time

Dataset Size	Single Row Insertion Time in Seconds	
	MySQL	HBase
30,000	38	5
60,000	68	8
90,000	95	9
1,30,000	134	10
2,00,000	202	16
2,50,000	255	18
5,00,000	523	39

Within the batch insertion experiment we find the number of inserts per second for different batch sizes. We calculated inserts per second by dividing total inserts with the total time taken in seconds. Fig. 5(a) illustrates that the number of inserts per second decreases as batch size increases in HBase while in MySQL it remains constant. On an average, number of inserts per



(a) Number of Inserts per Second in Batch



(b) Single Row Insertion Time in Seconds

Figure 5: Number of Inserts per Second in Batch and Single Row Insertion Time in Seconds

second in HBase is 17 times more in comparison to MySQL. The results are shown in Fig. 5(a) and Table 9. As can be seen from Fig. 4(b) and Fig. 5(a), the performance of HBase is better as compared to MySQL. For batch insertion MySQL uses InnoDB default buffer size of 8 MB to add batch records in it until buffer reaches a threshold value or commit occurs. On the other hand, HBase stores all its batch records in HBase write client buffer which is configured as 20 MB in HBase configuration file. We perform an experiment with the same configuration. Thus, we change InnoDB buffer size from 8 MB to 20 MB.

In HBase there is a lag in persisting the data stored on memstore to disk and it is by default asynchronous. On the other hand, MySQL persists data on disk and it is by default synchronous. To have the same configuration we change the durability of HBase to `FSYNC_WAL` in HBase configuration file. `FSYNC_WAL` writes the data to WAL synchronously and forces it to the disk. From the results it can be seen that time taken in HBase is 25 times lower in loading 5,00,000 records with different batch sizes as compared to MySQL. We believe the reason for this can be the difference in the way records are inserted in MySQL and HBase. In MySQL, executing an insert statement is a five step process. The batched insert statements in a buffer are first sent to the server, then parsed, then values are checked for uniqueness (intent hidden query), then data is inserted in actual table and finally data is inserted in index table. In HBase executing an insert statement is a two step process. The first step is writing the data to WAL then to the memstore and finally to the disk synchronously. Thus, the performance of HBase is better as compared to MySQL for batch insertion.

We also conduct a single row insertion experiment to examine which database can perform better for single row insertion. Fig. 5(b) and Table 10 reveal that the performance of HBase is better as compared to MySQL for all the datasets. The reason is same as batch insertion but here instead of sending records in a batch we are sending a single record in a single round trip. Fig. 5(b) reveals that when the dataset size increases then the difference between the time taken in loading real time data in MySQL and HBase also increases. We examine that the difference is 14 times lower in HBase as compared to MySQL. Hence, performance of HBase is better as compared to MySQL in loading different datasets with single record insertion.

8 Conclusion

In this paper, we present an implementation of α -miner algorithm in MySQL and HBase using SQL. Furthermore, we present the performance benchmarking and comparison of α -miner algorithm in MySQL and HBase. The α -miner implementation in MySQL is a one tier application which uses only standard SQL queries and advanced stored procedures. Similarly, implementation in HBase is done using Phoenix. We conclude that HBase on an average is 29 times faster than MySQL in loading large datasets. Based on experimental results, we conclude that HBase outperforms MySQL in loading real time data (event logs) by having 17 times more number of inserts per second.

We conclude the total time taken to read the data while execution of α -miner algorithm is 1.16 times lower in HBase as compared to MySQL. Similarly, for writing the data, time taken by HBase is 1.70 times lower as compared to MySQL. We conclude the total execution time of α -miner algorithm improves significantly in HBase as compared to MySQL by 1.46x order of magnitude. HBase outperforms MySQL in terms of the disk usage of tables. The disk space occupied by tables in HBase is 4.37 times lower as compared to MySQL. Thus, we conclude that HBase is more efficient than MySQL in terms of storing data and performing query. Using well known compression technique, HBase outperforms MySQL in disk usage as well as execution of α -miner algorithm.

References

- [1] Anand, K., Gupta, N., Sureka, A.: Utility-based control flow discovery from business process event logs. In: Proceedings of the 4th International Conference on Big Data Analytics - Volume 9498. pp. 69–83. BDA 2015 (2015)
- [2] Appuswamy, R., Gkantsidis, C., Narayanan, D., Hodson, O., Rowstron, A.: Scale-up vs Scale-out for Hadoop: Time to Rethink? (2013)
- [3] Chamberlin, D.D., Astrahan, M.M., Blasgen, M.W., Gray, J.N., King, W.F., Lindsay, B.G., Lorie, R., Mehl, J.W., Price, T.G., Putzolu, F., Selinger, P.G., Schkolnick, M., Slutz, D.R., Traiger, I.L., Wade, B.W., Yost, R.A.: A History and Evaluation of System R (1973)
- [4] C.Ordonez, P.Cereghini: Sqlem: Fast Clustering in Sql using the EM Algorithm. International Conference on Management of Data pp. 559–570 (2000)
- [5] Dimiduk, N., Khurana, A.: HBase In Action
- [6] D.J.Abadi, S.R.Madden, N.Hachem: Column-Stores vs. Row-Stores: How Different Are They Really? SIGMOID (2008)

- [7] D.P.Rana, N.J.Mistry, M.M.Raghuwanshi: Association Rule Mining Analyzation Using Column Oriented Database. *International Journal of Advanced Computer Research* 3(3), 88–93 (2013)
- [8] Fay Chang, Jeffrey Dean, S.G.W.C.H.D.A.W.M.B.T.C.A.F., Grube, R.E.: Bigtable: A Distributed Storage System for Structured Data. *OSDI'06: Seventh Symposium on Operating System Design and Implementation*
- [9] Finn, M.A.: Fighting impedance mismatch at the database level
- [10] Gupta, K., Sachdev, A., Sureka, A.: Pragamana: Performance comparison and programming alpha-miner algorithm in relational database query language and nosql column-oriented using apache phoenix. In: *Proceedings of the Eighth International C* Conference on Computer Science & Software Engineering*. pp. 113–118. C3S2E '15 (2008)
- [11] Gupta, M., Sureka, A.: Nirikshan: Mining bug report history for discovering process maps, inefficiencies and inconsistencies. In: *Proceedings of the 7th India Software Engineering Conference*. pp. 1:1–1:10. ISEC '14 (2014)
- [12] Joishi, J., Sureka, A.: Vishleshan: Performance comparison and programming process mining algorithms in graph-oriented and relational database query languages. In: *Proceedings of the 19th International Database Engineering & Applications Symposium*. pp. 192–197. IDEAS '15 (2014)
- [13] K-U.Sattler, O.Dunemann: Sql Database Primitives for Decision Tree Classifiers. *Conference on Information and Knowledge Management* pp. 379–386 (2001)
- [14] Kundra, D., Juneja, P., Sureka, A.: Vidushi: Parallel implementation of alpha miner algorithm and performance analysis on cpu and gpu architecture. In: *International Conference on Business Process Management*. pp. 230–241. Springer (2015)
- [15] Li, F., M.Tamer, Chen, G., Ooi, B.C.: R-Store: A Scalable Distributed System for Supporting Real-Time Analytics
- [16] L.Suresh, Simha, J., Velur, R.: Implementing k-means Algorithm using Row Store and Column Store Databases-A Case Study. *International Journal of Recent Trends in Engineering* 4(2) (2009)
- [17] Lübcke, A., Saake, G.: Workload Representation across Different Storage Architectures for Relational Dbms (2012)
- [18] Plattner, H.: A Common Database Approach for Oltp and Olap using an In-Memory Column Database. *ACM SIGMOD International Conference on Management of data* (2009)

- [19] R.G.Mehta, N.J.Mistry, M.Raghuvanshi: Impact of Column-Oriented Databases on Data Mining Algorithms. International Journal of Advanced Research in Computer and Communication Engineering pp. 2503–2507 (2013)
- [20] Russell, N.C.: Foundation of Process-Aware Information Systems
- [21] Sachdev, A., Gupta, K., Sureka, A.: Khanan: Performance comparison and programming alpha-miner algorithm in column-oriented and relational database query languages. In: Proceedings of the 4th International Conference on Big Data Analytics - Volume 9498. pp. 165–180. BDA 2015 (2015)
- [22] Shang, X., Sattler, K., Geist, I.: Efficient Frequent Pattern Mining in Relational Databases (2004)
- [23] Sharma, V., Dave, M.: Sql and NoSQL Database. International Journal of Advanced Research in Computer Science and Software Engineering (2012)
- [24] S.Weerapong, P.Porouhan, W.Premchaiswadi: Process Mining Using α -Algorithm as a Tool. IEEE (2012)
- [25] Vora, M.: Hadoop-Hbase for Large-Scale Data. Computer Science and Network Technology (ICCSNT), 2011 International Conference on (2011)
- [26] W.V.D.Aalst: Process Mining: Overview and Opportunities. ACM (2012)
- [27] Zhang, C., Sterck, H.D.: Hbasesi: Multi-Row Distributed Transactions with Global Strong Snapshot Isolation on Clouds. Scientific International Journal for Parallel and Distributed Computing (2011)